

Applicant : Norman Margolus et al.
Serial No. : 09/785,535
Filed : February 16, 2001
Page : 26 of 32

Attorney's Docket No.: 11656-002001

REMARKS

Appreciation is expressed for the personal interview accorded the undersigned and the inventor, Norman Margolus, in July 2005. The examiner and his SPE, Shahid Alam, attended the interview.

An amendment to the next to the last line of claim 1, to reinstate the use of "physical locations" in place of "physical storage node" was presented at the interview. The examiner had objected to "physical storage node" as lacking antecedent basis in the specification. (Applicant does not agree with the examiner's objection, and reserves the right to pursue an independent claim using the "physical storage node" language in a future continuation.)

The examiner's rejection of claim 1 under 35 USC 112, paragraph 1, for failure to provide a written description was discussed. The examiner objects to the "relying" element of the claim, which read:

relying on the data item already stored in the data repository for storage of the second data item rather than storing a separate copy of the second data item

Discussion made it clear that the examiner would prefer that the element be expressed as a more concrete action. It was proposed to the examiner that the element be rewritten to refer to the storing (or not storing) actions that take place depending on the outcome of the comparing element, and the examiner agreed that this would overcome the rejection. (Applicant does not agree with the examiner's rejection of the "relying" element, and reserves the right to pursue a claim using the relying element in a future continuation.)

Claim 1 has now been amended as discussed at the interview. The following element replaces the "relying" element:

storing the second data item in the data repository if comparing establishes that a data item identical to the second data item is not already stored in the data repository, and not storing the second data item in the data repository if comparing establishes that a data item identical to the second data item is already stored in the data repository

Applicant : Norman Margolus et al.
Serial No. : 09/785,535
Filed : February 16, 2001
Page : 27 of 32

Attorney's Docket No.: 11656-002001

For the new language to be consistent with the remainder of the claim, the claim has been further amended in two places so that the second data item may be either identical or not identical to data items already stored in the data repository. First, the "having a second program" step has been changed to remove the requirement that the second data item be identical:

having a second client program initiate a process for depositing a second data item in the data repository, [wherein the second data item is identical to the data item stored by the first client program,]

Second, the comparing step has been changed to have it determine "whether" the second data item was identical (rather than "that" is was identical):

comparing the digital fingerprint from the second data item to digital fingerprints for data items already stored in the data repository, and determining from the comparing of digital fingerprints, without comparing the entire contents of the second data item to the entire contents of a data item already stored, [that] whether a data item identical to the second data item is already stored in the data repository

The examiner raised a further objection to the language of claim 1 at the interview, one not made in the office action. The examiner objected to the use of "adapted to" in the first of the two "wherein" clauses at the end of the claim. It was proposed that "adapted" be replaced by "designed", and the examiner indicated that the change would overcome his objection. (Applicant does not agree with the examiner's objection to the use of "adapted to", and reserves the right to pursue a claim using that expression in a future continuation.)

Minor further amendments have been made to the dependent claims for consistency with claim 1.

Three new dependent claims (187-189) have been added. They are supported by disclosure in the original application. Support for claim 187 is found at paragraph 63 of the published application, which includes the sentence:

Applicant : Norman Margolus et al.
Serial No. : 09/785,535
Filed : February 16, 2001
Page : 28 of 32

Attorney's Docket No.: 11656-002001

All clients which read using named objects (such as 3d and 10) that are associated with the same dataname 3a actually share access to a single repository data-item 3.

Support for claim 188 is found at paragraph 63, in the following sentence:

A read client 5 (Fig. 3) associated with client 1 can use the access-authorization credential 3b that was generated in the deposit transaction to subsequently read data-item 3 indirectly by reference to named-object 3d, but no client can directly read data-item 3.

Support for claim 189 is found in paragraph 113, immediately after the reference to Fig. 9:

When write client 56 associates a new data-item 62 with named object 58, the reference count of the previous data-item 60 associated with named object 58 may go to zero. This means that data-item 60 is unreferenced, and it may be deleted and its storage reclaimed.

The examiner's rejection of claim 1 under 35 USC 103 as being unpatentable over Shnelvar combined with Waters was addressed at the interview.

It was pointed out how claim 1 differs from the prior art in at least two important respects.

First, contrary to what the examiner has concluded in the office action, Shnelvar does not teach using the pseudorandom distribution of the digital fingerprints (hash values) to introduce pseudorandomness into the physical storage location of the data items. Shnelvar uses the hash value to look up a pointer 66 (in the MDC Table, Fig. 3) that prescribes the physical location of the data unit (the "RAU" at which the "data unit" is stored).¹ By using the hash value to look up a pointer, and using the pointer to prescribe the physical location, Shnelvar teaches away from

¹If there is a "collision" – i.e., a new data item has the same hash value as an existing data item – Shnelvar uses a different location to store the data item, but that, too, teaches away from the invention, which does not store the second data item when its digital fingerprint is the same as a fingerprint of a data item already stored in the repository.

Applicant : Norman Margolus et al.
Serial No. : 09/785,535
Filed : February 16, 2001
Page : 29 of 32

Attorney's Docket No.: 11656-002001

using the pseudorandomness of the hash value to introduce pseudorandomness into the physical locations. The pseudorandomness of the hash value is lost by using the table of pointers.

This difference from Shnelvar can be understood by considering an example explained at the interview. Assume that you have two different data items that are stored in a data repository in succession, i.e., one after the other. Because the data items are different, their digital fingerprints are pseudorandomly different. With the invention, the pseudorandomness of the digital fingerprints is introduced into the physical locations at which the data items are stored, with the result that the two data items would be stored at pseudorandomly different physical locations. But in Shnelvar the data items would very likely be stored at adjacent physical locations. The fact that they are stored in succession would likely mean that they would be adjacent each other in the pointer table. Being adjacent in the pointer table would mean that are likely to have adjacent pointer values, and thus adjacent physical addresses (which the pointer values prescribe).

Thus, if a data file consisting of multiple blocks (data items) was being stored one block after another in succession, the blocks would likely be stored at pseudorandomly different physical locations with the invention but at adjacent locations with Shnelvar.

The examiner agreed at the interview that this teaching in Shnelvar was different from the invention (but the examiner did reserve the right to review Shnelvar again to confirm that there was nothing elsewhere in the reference that taught the claimed feature of introducing pseudorandomness into the physical storage locations).

The second difference from claim 1 that was discussed at the interview had to do with the teaching of Waters. The examiner relies on Waters as teaching another limitation of claim 1, a limitation that he concedes is not taught by Shnelvar:

determining from the comparing of digital fingerprints, without comparing the entire contents of the second data item to the entire contents of a data item already stored, whether a data item identical to the second data item is already stored in the data repository.

Applicant : Norman Margolus et al
Serial No. : 09/785,535
Filed : February 16, 2001
Page : 30 of 32

Attorney's Docket No.: 11656-002001

Shnelvar teaches away from the invention in this respect. If a comparison of digital fingerprints suggests that data items might be identical, Shnelvar teaches that the entire contents of the data items must be compared – precisely the opposite of claim 1.

The examiner turns to Waters for disclosure of this limitation of the claim. But, as explained at the interview, Waters, too, fails to teach the limitation.

Waters is concerned with retrieving content from a web server. When a browser requests content using a URI, Waters forms the hash of the URI, uses the hash to lookup a pointer, and then retrieves whatever it finds at the physical location specified by the pointer.² Waters admits that this will result in the wrong content being delivered, as he expects that “hash collisions” will occur, i.e., the same digital fingerprint will be produced by different URIs. But Waters teaches that, at least in the web server application described, such errors can be tolerated, and there are performance advantages to delivering content even with a likelihood that, on occasion, what is delivered is wrong.

Thus, Waters does teach making a decision to retrieve data based on identity of digital fingerprints, without comparing the entire contents of the URIs, but Waters fails to teach the remainder of the limitation. We quote the limitation once again, and underscore those portions that Waters fails to teach:

determining from the comparing of digital fingerprints, without comparing the entire contents of the second data item to the entire contents of a data item already stored, whether a data item identical to the second data item is already stored in the data repository.

Waters does not teach comparing digital fingerprints to determine whether data items are identical. Quite to the contrary, Waters teaches that comparing digital fingerprints cannot determine whether data items are identical and that failing to know whether the data items are identical does not matter.³

²Another difference from Waters is that Waters forms the digital fingerprint not of the content of a data item stored in a repository but of a name (URI) for that data item.

³We speak of Waters determining whether “data items” are identical, but as we noted earlier, Waters is actually determining whether the name of a data item is identical.

Applicant : Norman Margolus et al.
Serial No. : 09/785,535
Filed : February 16, 2001
Page : 31 of 32

Attorney's Docket No.: 11656-002001

The examiner's combination, even if it were made, would thus not reach the invention, for the one thing that both Shnelvar and Waters have in common is that they both fail to teach that one can determine whether data items are identical by comparing their digital fingerprints without comparing the entire contents of the data items. Any combination of the two references would thus also fail to teach this aspect of claim 1.

But actually Shnelvar and Waters differ so dramatically in what they teach that it is not reasonable for the examiner to assume that Waters would be used to modify the teaching of Shnelvar. Shnelvar insists that one must compare the entire contents of the data items, and cannot rely on identity of digital fingerprints. Waters insists that the exact opposite is the better procedure -- that the entire contents (of the URIs) do not have to be compared. Waters reasons that in the web server application with which it is concerned, it is better on balance to occasionally make a mistake and deliver the wrong content.

After having these points explained at the interview, the examiner indicated in the interview summary record that "claim 1 appears to define over the prior art discussed [i.e., Shnelvar and Waters]".

The remaining claims, which were not discussed at the interview,⁴ are all properly dependent on claim 1, and thus allowable therewith. Each of the dependent claims adds one or more further limitations that enhance patentability, but those limitations are not presently relied upon. For that reason, and not because applicants agree with the examiner, no rebuttal is offered to the examiner's reasons for rejecting the dependent claims.

Allowance of the application is requested.

⁴Dependent claim 185 was briefly discussed, and agreement was reached that the amendment made herein would overcome the rejection under 35 USC 112, first paragraph.

Applicant : Norman Margolus et al.
Serial No. : 09/785,535
Filed : February 16, 2001
Page : 32 of 32

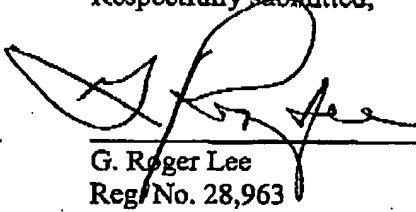
Attorney's Docket No.: 11656-002001

Please apply any other charges or credits to deposit account 06-1050.

Respectfully submitted,

Date:

8/25/05



G. Roger Lee
Reg/No. 28,963

Fish & Richardson P.C.
225 Franklin Street
Boston, MA 02110
Telephone: (617) 542-5070
Facsimile: (617) 542-8906

21155228.doc